

ProLiant Essentials Workload Management Pack

Featuring Resource Partitioning Manager

best practice



Executive summary.....	3
Introduction.....	3
Simplify management.....	3
Increase utilization.....	5
Cluster "landing zones"	5
Best practices	6
General application notes.....	6
Performance impact of using RPM	6
Caveats	7
Privately managed resources	7
Single instance applications.....	7
General caveats.....	7
Databases	7
Oracle.....	7
Considerations.....	7
Microsoft SQL2000 database server	8
Messaging applications.....	8
Microsoft Exchange 2000.....	8
Test environment.....	8
Considerations.....	8
Lotus Domino	9
Considerations.....	9
Limitations	9
Enterprise applications	10
mySAP.com.....	10
Example 1	10
Example 2	11
Example 3	11
PeopleSoft.....	12
Test environment.....	12

Considerations.....	12
Web servers	12
Java-based	12
Apache.....	12
Microsoft IIS	12
Infrastructure	13
Microsoft SharePoint Portal Server.....	13
Testing environment	13
Considerations	13
Citrix Metaframe and Microsoft Terminal Services.....	13
Custom developed applications	14
File/print services.....	14
Clustering	14
Microsoft Cluster Server	14
Considerations	14
Limitations	15
Summary	15
Troubleshooting.....	15
Upgrading from the Resource Partitioning Manager Preview Version.....	15
Hardware configuration changes	15
GUI synchronization.....	15
Terminal Services	16
Running multiple instances of the same application	16
Application Active Resource Partition	16
RPM causes applications to run incorrectly	16
RPM causes other applications to fail or receive access violations	17
RPM does not allow processes to run.....	17
Additional information	17
For more information.....	18
Call to action	18



Executive summary

HP ProLiant Essentials Workload Management Pack featuring Resource Partitioning Manager (RPM) increases the stability and availability of applications under Windows 2000 to allow customers to confidently deploy multiple applications on a single server. RPM builds upon the embedded job object technology in Windows 2000 to provide a quick and easy way to manage processor and memory resources dynamically.

This white paper will assist you in understanding what applications benefit from the Workload Management Pack and will provide basic guidelines on how to set up those applications in a managed environment.

Introduction

Delivering a truly adaptive IT environment calls for an IT partner with the vision to imagine the future of IT, the ability to create it, and the commitment, expertise, and resources to deliver it. Adaptive Infrastructure represents HP's vision to deliver more flexible, intelligent, automated IT infrastructures that adapt to the unique demands of your business. Compaq is delivering on its Adaptive Infrastructure blueprint through its roadmap of innovative ProLiant technologies including the ProLiant Essentials Workload Management Pack featuring Resource Partitioning Manager.

For customers requiring dynamic resource control in a "scale-up" model, the Workload Management Pack provides easier management of complex environments, improving overall server utilization and enabling Windows 2000 customers for the first time to confidently deploy multiple applications on a single multiprocessor ProLiant Server. Now ready for deployment in "production" environments, the Workload Management Pack, featuring HP Resource Partitioning Manager, increases server utilization by giving administrators control over the size and physical location of system resources available to individual applications, services and other processes.

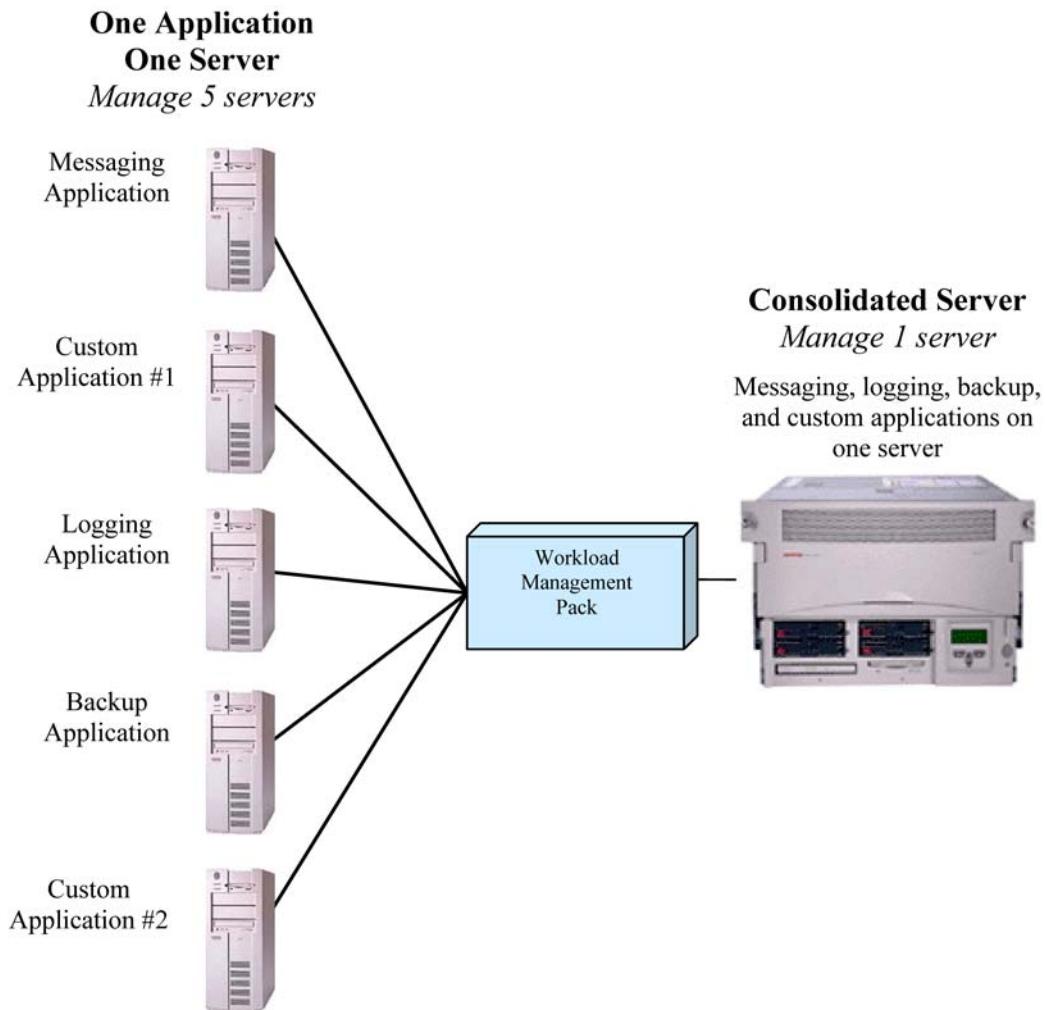
Resource Partitioning Manager (RPM) is an easy-to-use, GUI-enabled tool that extends the Windows 2000 operating system to give IT administrators the power to dynamically optimize their ProLiant servers. With RPM, you can easily manage your complex Windows 2000 environments and confidently deploy multiple applications on a ProLiant server, improving overall server utilization. RPM users can create multiple resource partitions, limit each partition to specific resource quantities, and establish rules that allow for the dynamic reallocation of processors and memory. The ability to assign flexible resource boundaries to partitions enables customers to confidently deploy multiple applications on a ProLiant server.

Environments that benefit most from using the Workload Management Pack featuring Resource Partitioning Manager are those that include applications currently running on individual servers but utilizing few system resources. For example, because many applications use system resources in a conflicting manner each application is isolated on separate machines. Consolidating these applications allows IT personnel to reduce administration costs, reduce overall software licensing fees, and reduce overall hardware costs.

Simplify management

Resource Partitioning Manager provides a flexible way to make server consolidation a reality by reducing costs associated with the management of one server for each application (see Figure 1). Using the partition configuration wizard, resource partitions can quickly be created to reserve resources for a group of processes, shielding applications from other resource consuming processes running on the server.

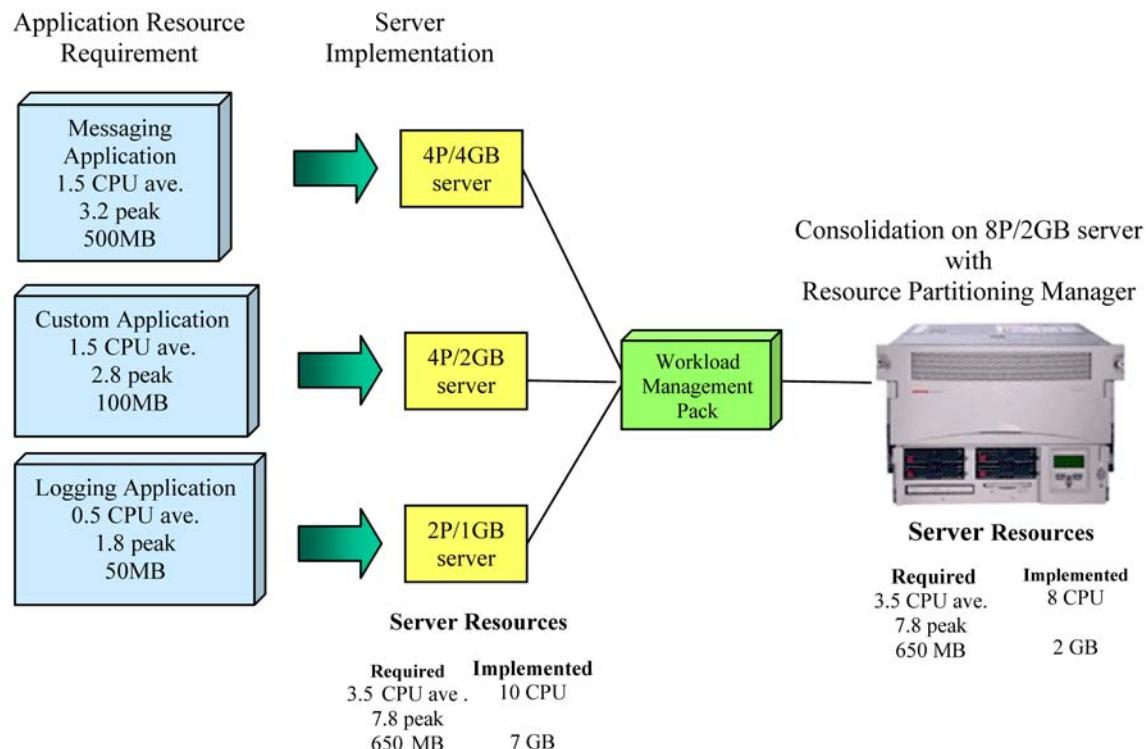
Figure 1. Server consolidation using Workload Management Pack



Increase utilization

The ability to respond quickly to increasing processing demands has become critical in current IT infrastructures. As a result, server environments contain numerous servers that are underutilized. These servers hold valuable resources idle in anticipation of an unforeseen demand. With RPM, overall headroom can be reduced without compromising the amount of resources available to each application (see Figure 2). Applications maintain access to a comparable amount of resources that can automatically be reallocated to those processes experiencing peak loads.

Figure 2. Consolidating resources onto one server versus multiple servers



Cluster “landing zones”

Many clustered Windows 2000 environments are still configured with one active node that runs the high availability application and a second node, configured as a passive node, ready to accept the high availability application in the event of a failure. This configuration ensures that the high availability application has access to the resources it needs during fail-over if other applications are also running on the server. With RPM, users can create “landing zones,” making it possible to convert passive nodes to active nodes. “Landing zones” ensure that system resources are available once the fail-over takes place, without requiring that an entire server sit idle in reserve. Now you can control the workload of a cluster not only during normal operations, but also during fail-over.

Best practices

Workload Management, or Resource Partitioning, brings many of the benefits associated with hardware and logical partitioning. It provides the resource control needed to run multiple applications on a single server without the added capital cost of a server capable of hardware partitioning. Workload management also reduces the complexity of logical partitioning by managing only specific resources such as memory and processors, allocating these resources to specific applications. However, partitioned systems can still take advantage of Workload Management and Resource Partitioning within a hardware or logical partition.

When determining what types of applications benefit the most with RPM, consider what purpose is being met:

- Heterogeneous application consolidation – This involves consolidating many different applications, typically custom or in-house developed applications that today, run on their own individual servers, onto fewer larger servers.
- Multiple instance application consolidation – Many applications allow multiple instances of the application to run on a single server. This involves consolidating those multiple instances currently running on dedicated machines, onto a larger server. Until now there was no way to control how the resources were distributed among the individual instances and so these had to be deployed on individual servers.
- Other types of server consolidation – RPM supports MSCS clustered environments and also allows for “landing zones”. “Landing zones” are partitions that reserve resources for applications that fail-over from the primary server.
- Resource control of misbehaved applications – RPM allows control over CPU and memory resources and can limit applications that leak memory or consume CPU time without regard for the rest of the server. Often these applications will bring the entire system down requiring a system reboot.

General application notes

This section contains descriptions and test scenarios of multiple applications that have been tested using Resource Partitioning Manager Requirements. In general, any application that runs in a Windows 2000 environment (Server, Advanced Server, or Datacenter Server) can be controlled using RPM. Because RPM controls resources within a single operating system, consolidated applications under RPM all run on the same operating system instance.

Note:

Because HP cannot test every combination of every application with RPM, it is recommended that users deploying RPM conduct testing of their specific application combinations prior to placing RPM and the consolidated application environment into production use.

Global Services offers packaged and custom services to support your testing and server consolidation requirement.

Performance impact of using RPM

HP has conducted performance testing of individual, commercially available applications running with RPM and has found significantly less than 1% performance impact when compared to running the same individual application on the same server. In some cases, use of the advanced settings within RPM can decrease context switching and actually improve application performance.

Caveats

There are certain caveats to consider when running certain applications in an RPM environment.

Privately managed resources

Some applications, such as Microsoft SQL Server, already contain resource management functionality similar to that found in RPM. With applications like Microsoft SQL Server 2000, the built-in resource management functionality should be used to control resources for that application. RPM can then be used on the same server to ensure that other running applications do not encroach upon the required resources for privately managed applications like SQL Server.

Single instance applications

Some products, such as Microsoft Exchange Server, Microsoft Internet Information Server (IIS), and Terminal Server do not allow multiple instances to run under a single operating system instance. Use of RPM will not overcome these product-specific limitations. For example, you may run Microsoft Exchange Server from within a Resource Partition on the same computer running IIS from within another Resource Partition but you may not simultaneously run two copies of Microsoft Exchange Server from within two Resource Partitions on the same OS instance.

General caveats

- It is not recommended that you run RPM remotely on a system using Microsoft Terminal Server.
- See the Troubleshooting section of this paper for more information on running multiple instances of the same application.

Databases

Oracle and Microsoft SQL2000 Database Server have both been tested while running RPM.

Oracle

Customers choosing to deploy HP's Workload Management Pack featuring Resource Partitioning Manager with Oracle databases and applications on Windows 2000 Server, Advanced Server, or Datacenter Server will be supported by Oracle under their normal Oracle Service agreements. Oracle and Resource Partitioning Manager version 1.1 have been tested and proven to run without conflict on ProLiant Servers.

The testing for Oracle and RPM was conducted on a ProLiant DL580 with four gigabytes of memory and four 700 MHz processors. During the testing, the number of users, processors, and memory was varied to get the most thorough results. Both single and double instance tests were run on the same box using RPM to manage the memory and processor resources for the Oracle instances.

The testing indicated no significant difference between the aggregate performance of 2 two processor servers running with one instance of Oracle each versus running 1 four-processor server with two instances of Oracle running in separate Resource Partitions. RPM was able to add and remove processors as the workloads needed the extra capacity. All testing was performed using Windows 2000 Advanced Server and Oracle 8i, version 8.1.7.

Considerations

Server consolidation of multiple Oracle database servers is possible under RPM. Multiple instances of Oracle can also run together on Windows 2000 without a workload management tool. However, RPM makes the management of the consolidated environment easier through its ability to easily assign resources. Oracle manages memory by itself and will not grab additional available memory if it does not need it. It only uses the memory that was assigned to it through the parameter file. Because of this characteristic, memory rules would not typically be invoked for an Oracle partition.

Microsoft SQL2000 database server

As previously mentioned, HP and Microsoft recommend that customers use the built-in resource management functionality within the SQL administrative tools to control resources for Microsoft SQL2000 Database Server. However, multiple instances of Microsoft SQL2000 Database Server can be run on a single server. The same set of administrative tools can be used to control the resources available to these multiple instances of the SQL database. RPM can be used on this same server to ensure that other applications running do not encroach upon the required resources for a single instance or multiple instance deployment of Microsoft SQL2000 Database Server.

Messaging applications

The messaging applications tested were Microsoft Exchange 2000 and Lotus Domino.

Microsoft Exchange 2000

Microsoft Exchange 2000 is an example of an application that runs as a single instance on a server. Because of this, it is not possible to run multiple Exchange 2000 servers on a single operating system instance. However, as customers are migrating from an Exchange 5.5 to an Exchange 2000 environment, they are increasingly putting more mail users on each server and combining other applications on the Exchange 2000 server to improve the overall functionality and robustness of their Exchange environment. Because of this shift in server configuration, initial Exchange 2000 testing with RPM was performed to determine how Exchange 2000 could benefit from Resource Partitioning.

Test environment

The test environment consisted of a ProLiant 8500 with two gigabytes of memory and eight 700 MHz CPUs running Windows 2000 Datacenter and Exchange 2000. Microsoft's Loadsim program was used with the MMB2 script to generate 5000 simulated Microsoft MAPI mail users. A simple test was conducted with a resource partition to control the STORE.EXE process.

The STORE.EXE process is a single process that controls access to Exchange 2000 information stores. When multiple virtual servers and multiple information stores are configured on a single physical server, all information store databases are accessed through a single instance of the STORE.EXE process. Because this process is typically the bottleneck in an Exchange 2000 benchmark, STORE.EXE was chosen as the process to test with RPM. The initial step in testing RPM with Exchange 2000 was to see if STORE.EXE could be controlled by RPM.

The resource partition was configured to allow the STORE.EXE process to run on two CPUs and the Schedule Class was changed from Normal to Above Normal. A rule was enabled to add additional CPUs to the resource partition when total CPU utilization consumed by the STORE.EXE process increased above 70% for 30 seconds. The system dynamically added an additional four CPUs to the resource partition for a total of six CPUs allocated to the STORE.EXE process for the duration of the test. Using RPM, we were able to control the STORE.EXE process and allocate additional CPU resources as required.

Context switching appears to be reduced by setting the schedule class for STORE.EXE from normal to above normal using the advanced tab screen within RPM. With RPM schedule class set to above normal the context switching measured an average of 7800 context switches per second. Without RPM, the context switches per second measured an average of 9200. This resulted in a decrease in context switching of approximately 1400 context switches per second.

Considerations

Testing of memory allocation to the store process could also result in some benefit surrounding recent problems with cluster fail-over and memory fragmentation associated with the STORE.EXE process during a cluster fail-over. This problem currently limits the recommended number of concurrent users in

an Active/Active cluster configuration to 1900 concurrent users per Exchange 2000 Virtual server SP2 with proactive monitoring of the cluster to ensure that the CPU does not exceed 40 percent (load generated from users) loading. For example, a "Landing Zone" resource partition could be used to allocate memory to the STORE.EXE process to help ensure that the information store process has sufficient memory to operate in the fail-over situation.

Additional testing is currently underway to validate advanced configurations and characterize performance of consolidated Exchange 2000 environments.

Lotus Domino

Lotus Domino is an application that will run multiple Domino instances on a single server. Lotus Domino is, therefore, an excellent candidate for consolidation. To configure this environment, administrators must create multiple partition installs of Domino and then assign the different executables separately to individual resource partitions with RPM. Basic information on the configuration of single and multiple instances of Lotus Domino Server can be found at www.10.lotus.com/ldd/notesua.nsg.

Considerations

Because there are over 60 executables on the Lotus Domino server, administrators must determine which executable they want to assign to the different RPM partitions depending upon the type of work the Domino server will be doing and their specific environment. Some of the primary executables to capture in individual partitions are:

- NSERVER.EXE – the main Domino server executable
- NREPLICA.EXE – the replication server task if there are heavy replication schedules in place
- NCLREPL.EXE – the cluster replication server task
- NROUTER.EXE – the server task that routes mail
- NDECS.EXE – the server task that allows the importing of other file types from other platforms

For a complete list and description of all Domino server tasks refer to the Lotus Domino "yellow book" titled *Administering the Domino System Volume 2* on page 1099. You can access this document through the Internet using the following link:

www10.lotus.com/ldd/notesua.nsf/0b345eb9d127270b8525665d006bc355/69f5e7de5964365e8525673e00491995?OpenDocument&Highlight=2,Domino,Administration.

Limitations

Because IBM recommends that many of the Lotus applications (Domino Mail server, Domino Fax server, etc.) run on separate boxes for performance reasons, RPM may or may not be the tool to satisfy the customer's Domino consolidation needs. RPM should be tested in a development/QA system configuration with the customer's specific Lotus Domino environment to verify that satisfactory performance and response times are achieved prior to placing the Lotus Domino/RPM configuration into production use.

HP tested RPM with Domino R5.06a Mail server and a Benchmark load of 4000 users on a HP ProLiant 7000, with four 500 MHz processors, 280-GB disk storage, and three GB of memory. The Benchmark test bed included one parent driver and four child drivers with 1000 users each.

When testing RPM on a Domino server, the nserver.exe task was assigned to two of four processors with a Scheduling Class of 7 and included a rule to add the remaining two processors if the CPU reached a capacity of 50% or more for one minute. It was determined that RPM would benefit users who plan to deploy more than just a single Domino application on one physical machine. By doing this, the user could assign individual processors to specific applications.

Enterprise applications

Enterprise applications tested with RPM include mySAP.com and PeopleSoft.

mySAP.com

SAP AG announced in 2001 that companies can now achieve true e-business flexibility and lower cost of ownership with the mySAP.com® e-business platform. Companies using mySAP.com now have the option to deploy several independent mySAP.com components not only on one server system but also on one database. RPM is the perfect tool to support the consolidation of several mySAP.com components on a single server system.

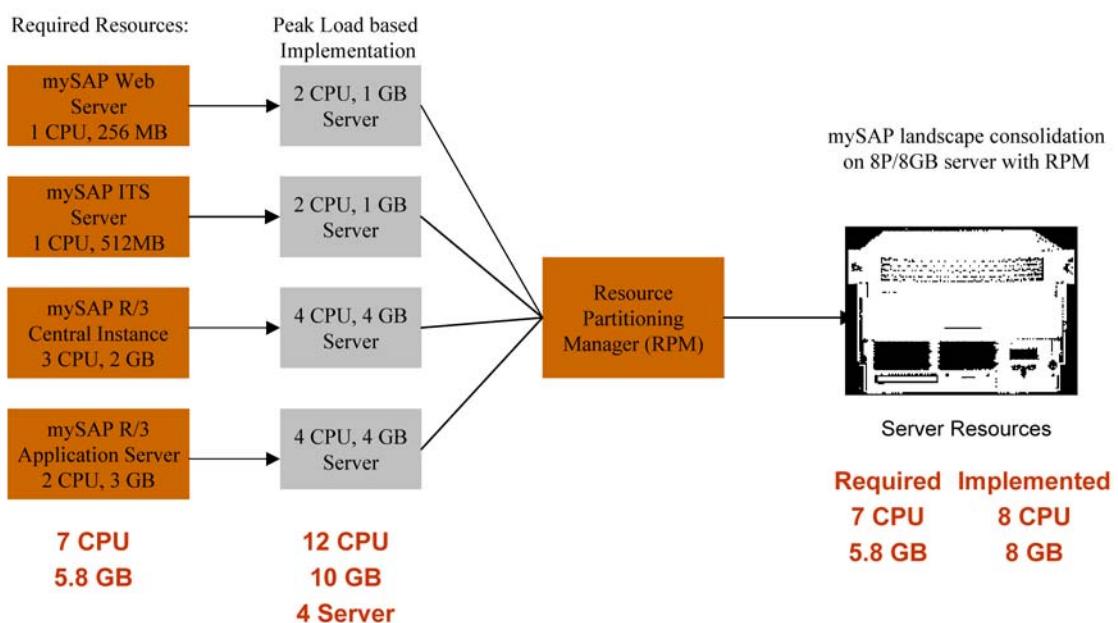
The mySAP.com components like SAP R/3, Workplace or CRM using the Windows 2000 job object APIs can be easily deployed with RPM. The RPM job object capturing feature allows for the capturing of running SAP objects and uses these captured objects for the creation of resource partitions. Once a mySAP.com resource partition was created, rules can be assigned to dynamically assign and revoke CPU and memory resources. The RPM dynamic rules engine gives you the flexibility you need when you deploy multiple mySAP.com components on a single server.

Deploying multiple mySAP.com components in one server system can simplify administration efforts and lead to reduced maintenance and operating costs.

Example 1

When consolidating multiple mySAP.com components on a single server system using RPM, no special installation templates from SAP are needed. In this example, SAP system transports, system copies, and the individual system management are possible because each component uses its own physical database file (see Figure 3). The consolidation of mySAP.com systems and components will lower your administrative burden, hardware costs, and maintenance costs.

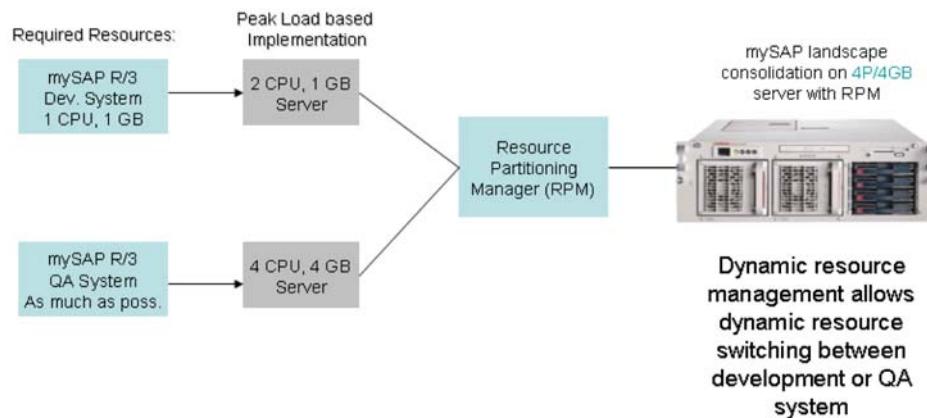
Figure 3: Consolidating multiple mySAP.com components onto one server using RPM.



Example 2

Use the following example when consolidating a mySAP.com development and quality system on one server system.

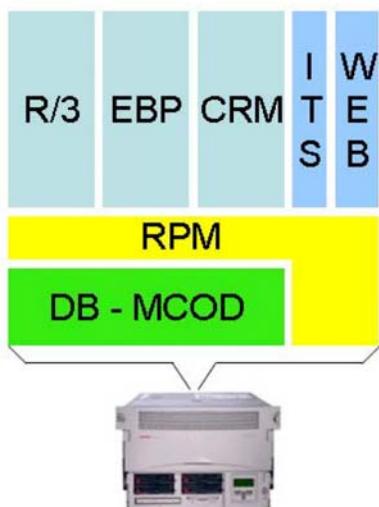
Figure 4: Consolidating a mySAP.com development and quality system onto one server.



Example 3

This example shows RPM combined with MCOD (Multiple Components – One Database). RPM supports the SAP consolidation solution with MCOD in dynamically balancing the available CPU resources between the different mySAP.com components. In addition to the reduced database, operating and server system maintenance efforts, consistent and reliable backup and restoration of business data is greatly eased since dependencies between different databases no longer occur.

Figure 5: Combining RPM with MCOD.



The main targets for MCOD mySAP.com installations are:

- Development systems
- Quality assurance systems
- Training systems
- Small to mid-range production systems

Based on a survey SAP did in 2001 and the figures they published¹, companies can save the following costs by consolidating their mySAP.com components:

- ~10% savings of disk space
- ~30% savings of hardware backup costs (tapes, tape drives, and disk systems).
- ~40% operating cost savings of such as managing backups or high-availability solutions.

¹ Source: SAP white paper "Multiple Components in One Database", Version 1.1, November 2001.

The savings mean reduced maintenance and operating costs overall; as well as, a streamlined process for backup and high-availability solutions and an efficient use of available resources.

PeopleSoft

PeopleSoft testing was conducted to demonstrate that three-tier environments could be consolidated on a single 8-way server and to gain a quick estimate of the performance of the system.

Test environment

PeopleSoft CRM 8.0 was used on a ProLiant DL760 with eight 550 MHz CPU's, four GB of memory on Windows 2000 Advanced Server. The environment included Microsoft SQL2000 server, the PeopleSoft Application, and Weblogic java based web servers from BEA. Partitions were created for PeopleSoft and Java-based web servers based on Weblogic while Microsoft SQL was bound to two unused processors using the embedded resource tools in SQL. The demo was able to show acceptable support for up to 250 users per CPU used in the environment, which compares to results from a 30,000 user benchmark

(See <ftp://ftp.compaq.com/pub/products/servers/benchmarks/46337.pdf>)
that delivered 203 users per CPU used in the solution.

Considerations

PeopleSoft 8.0 architecture lends itself to the use of RPM in achieving server consolidation. This consolidation can occur across or in multiple tiers of a large installation in addition to being able to collapse all tiers in a smaller installation.

Web servers

Web servers addressed in this paper include Java-based web servers, Apache, and Microsoft IIS.

Java-based

RPM is an excellent tool to help enable your Java-based web server consolidation plans. By running multiple instances of the Java-based web server, each within a resource partition, you can fully maximize the number of connections to your enterprise applications while maximizing the utilization of your hardware and minimizing your operational costs. Configuration information will vary from web server to web server depending upon on the specific web server vendor.

Apache

Although common on the Linux operating system, Apache HTTP server is also available for Windows 2000. Version 1.3.23 is the most current stable version available at the time of this paper and is recommended for use. Version 2.0 for Windows 2000 promises significant performance improvements and would be the most likely candidate for testing Apache on Windows 2000 performance versus other HTTP servers and platforms. Information and updates are available at <http://httpd.apache.org/>.

Because Apache HTTP server can run multiple instances on a single server, instances of the web server (each within a resource partition) can fully maximize the number of connections to your enterprise applications while maximizing the utilization of your hardware and minimizing your operational costs.

Microsoft IIS

Microsoft Internet Information Server (IIS) runs as a single instance of the service. You cannot run multiple IIS HTTP services on a single instance of the operating system. To capture the IIS HTTP service, start the service and assign the inetinfo.exe process to the IIS resource partition.

RPM can bring benefit to servers running IIS alongside other applications. See Example 3 of the mySAP.com section under Enterprise Applications.

Infrastructure

The infrastructure applications used in this paper are Microsoft SharePoint Portal Server, Citrix Mainframe and Microsoft Terminal services, and general custom developed applications.

Microsoft SharePoint Portal Server

Microsoft SharePoint Portal Server (SPS) cannot be described as either a monolithic application or as just a "simple web site" or portal. It is better described as a web application, where the application logic and intelligence is provided in the main underlying code. Other than the SPS Document Management service, there is no real application executable. SPS uses a number of Windows 2000 system services to accomplish its work. These revolve around a few main areas—the Microsoft Exchange Information Store service (used for the SPS database), the WWW Publishing service (web pages) and the MS Search service—along with the dependant services for these, such as IIS Admin, Protected Storage, RPC, and LSASS.

In an SPS environment, the following processes/services are potential targets for capture in resource partitions:

- STORE.EXE - SPS Web store service
- INETINFO.EXE - Web publishing service
- MSSEARCH.EXE - SPS search service

Testing environment

Testing was performed on a ProLiant DL580 configured with four 550 MHz (1MB) CPUs, two gigabytes of memory, SmartArray 5302, 18 disks (internal and outboard storage), and a single 10/100 NIC. This system has been in use for all SharePoint Portal Server (SPS) testing to date and is optimized for SPS performance (system/cache tuning, logical volume layout, file placement, etc.).

The operating system used was Windows 2000 Advanced Server patched to SP1 with additional recent "hot-fixes". Microsoft SharePoint Portal Server v1.0 was configured to represent a typical workgroup server. All significant data structures (for example, databases, indexes, logs, workspaces, etc.) were spread across a number of logical volumes (RAID arrays) to optimize I/O.

In multiple different configurations using RPM in the SPS environment, results showed no appreciable impact to performance using RPM.

Considerations

RPM does not impact the performance of SPS and can be used to consolidate servers running a single instance of SPS with other applications. Testing RPM proved it to be easy to use, totally stable and predictable, and provided good real-time monitoring information for the environment.

Citrix Metaframe and Microsoft Terminal Services

Both the Citrix and the Microsoft products are single instance applications/services that cannot run as multiple instances on a single operating system instance. However, RPM can be used to capture these individual services themselves. Unfortunately, RPM version 1.1 cannot automatically capture a daughter process spawned by these services. In order to capture these daughter processes, which is what is being delivered to the remote users, RPM must be manually configured to assign the processes to one or many resource partitions or set up in a preconfigured capture mode for each process. Development is underway to improve the functionality of RPM in a terminal server/Citrix environment.

Additionally, when viewing the server screen using Terminal Services or Citrix, the user is actually seeing a reproduction of the server screen, specific to that user's activities (as if he is the only user). Because of this effect, running RPM using Terminal Services can produce some confusion. More specifically, applications launched by RPM when using Terminal Services will not appear to have been started, because they will have been launched outside of the user's client space. Thus, they will not appear in the user's Terminal Services window, although they will appear on the actual server console.

Custom developed applications

Custom developed applications must be tested on an individual basis prior to being deployed on a consolidated platform with RPM. If custom developed applications currently run on Windows 2000, there is a high probability that they can be used with RPM. Custom applications benefit from RPM by being protected from other misbehaving applications that over consume CPU or memory. RPM can enable server consolidation of many custom developed applications.

File/print services

Windows 2000 file and print services can be captured and controlled using RPM. In Windows 2000 the SPOOLSV.EXE process is the print spooler and the DFSSVC.EXE process is the distributed file system service that controls the file system. Assign these processes to resource partitions to control the amount of CPU and memory these functions use.

Clustering

This section discusses the caveats of running RPM with Microsoft Cluster Server software.

Microsoft Cluster Server

HP Resource Partitioning Manager can be successfully used in a clustered environment without interfering in normal cluster operations. RPM provides "Landing Zones" for applications in a clustered environment. A "Landing Zone" is a Resource Partition that has been predefined for an application on the application's standby cluster node.

When using clustering with RPM, it is best to pair applications that do not use the same resources. For example, do not pair applications that use the same memory.

To set up a "Landing Zone" partition on the fail-over node in a 2-node cluster, set up a resource partition as you normally would on the first node, which will normally be running the application. On the second node, set up a similar resource partition but be sure to assign the fail-over processes to the partition using the "Add Process by Name" function. Once created, activate the "Landing Zone" partition on the second server and it will be configured to provide guaranteed resources to failed-over applications.

Considerations

There are two ways to implement a two node cluster under the Microsoft Cluster Service (MSCS)-either both nodes are running applications or one node is running applications and the other node is passive, waiting to take over the application(s). Datacenter allows up to four active nodes or $n+1$ active+passive.

1. Active - Passive configuration: If the active server is running one or more applications under RPM and the two nodes are identically equipped, then the partition definitions are identical on both servers. There is no need to configure the servers differently. If the nodes are not identical (a 4-way passive node and an 8-way active node) then the partition definitions may be different.
2. Active – Active configuration: If both nodes run different applications, (for example, a database on one node and messaging on another) and are identically equipped, then the definitions can be

identical. Each application can either own certain processors, be allowed to use other processors, or be denied a processor completely. If the nodes are not identical, then the definitions may be different.

Limitations

Most of the cluster applications do not have distinct process names for each instance. This means that for fail-over scenarios of servers with multiple instances of the same application, there is no way for RPM to distinguish each individual instance from another as they start back up. Since resource partitions must be predefined and activated as "landing zones" waiting for these fail-over processes to start on the server in response to a high-availability event, each partition that is designated for an application, should be configured so that any of the instances could eventually be active within that partition and should, therefore, be sized to accommodate this.

Summary

The ability to do more with less is central to the role of IT in today's economic environment. Today's business reality requires that IT projects have a short ROI and in many cases those projects include Server Consolidation. Workload Management, common on the mainframe, is now enabling application consolidation on industry standard platforms through the use of HP's Resource Partitioning Manager. The ProLiant Essentials Workload Management Pack featuring Resource Partitioning Manager is a direct response to upper management demands that resources be more fully utilized and is a key piece of any plan to consolidate industry standard servers and applications.

Troubleshooting

This section details known issues using the HP ProLiant Essentials Workload Management Pack and provides information about resolving them.

Upgrading from the Resource Partitioning Manager Preview Version

Due to significant enhancements to RPM features and formats, upgrades from the Resource Partitioning Manager Preview Version are not supported.

If you wish to install RPM on a computer that already has the Preview Version installed, you must uninstall the Preview Version by using Control Panel, Add/Remove Programs. When the program has successfully uninstalled, you can proceed with a normal installation of Resource Partitioning Manager version 1.1.

Hardware configuration changes

In the event of a significant change to the system configuration (particularly the addition or removal of processors or memory), Resource Partition configurations may be rendered invalid. In this case, Resource Partition configurations should be reviewed for any needed updates. RPM will display a warning box noting that a system configuration change has occurred, and it will try to highlight potential areas of concern.

GUI synchronization

If a single server is managed simultaneously by two Resource Partitioning Manager GUIs (one local and one remote), GUI synchronization issues may result. For example, if the GUI is left running on the server and an administrator uses a GUI remotely to delete a Resource Partition, that change will not automatically be reflected on the RPM GUI running on the server. If the GUI is closed and restarted, all changes will be updated.

In general, the best practice is to close the RPM GUI when it is not in use. It is essentially a configuration interface—the RPM service controls all resource partition management.

Terminal Services

Terminal Services is designed to emulate a mainframe client-server environment. When viewing the server screen using Terminal Services, the user is actually seeing a reproduction of the server screen, centric to that user's activities (as if they are the only user). Because of this effect, running RPM using Terminal Services can produce some confusion. More specifically, applications launched by RPM when running Terminal Services will not appear to have been started, because they will have been launched outside of the user's client space. Thus, they will not appear in the user's Terminal Services window, although they will appear on the actual server console.

Running multiple instances of the same application

While RPM supports running multiple instances of most applications, there are situations where you must capture multiple instances of the same program. When this occurs, make sure that the processes are captured by the correct Active Resource Partition.

The following shows three instances of an application and three Active Resource Partitions, each waiting to capture an instance:

Application Active Resource Partition

Application Instance # 1 Active Resource Partition # 1 (waiting for application to start)

Application Instance # 2 Active Resource Partition # 2 (waiting for application to start)

Application Instance # 3 Active Resource Partition # 3 (waiting for application to start)

In situations where each Resource Partition is configured differently, start each Resource Partition manually; then, manually start the application instance to be captured. Continue the sequence until all of the instances have been started (activate a resource partition, start an application instance, activate a resource partition, start an application instance, and so on).

NOTE:

This scenario only applies if an instance MUST be captured by a specific Resource Partition.

RPM causes applications to run incorrectly

Cause: When RPM manages processes and applies certain rules, some applications will behave differently. In some cases, applications will not work within those rules or applied constraints. Memory limits for process groups are stringent and applications are not allowed to break those limits.

Problem 1 - Memory Constrained - An application can fail if RPM limits its partition to an amount of memory below the application requirement. Solve this by increasing the memory limit.

Problem 2 – Process Constrained - If the partition containing the application that does not work has set an Active Process Limit (Default = 0 = no limit) and that application creates other processes, the other processes may fail because the Active Process Limit was exceeded. The processes that continue running may not work properly because other processes failed.

Problem 3 – Per Process Memory Limit Mismatch - When a physical memory or virtual memory limit is in effect for a partition, that limit is applied to each process in the group. A limit that is suitable for one process could be inadequate for another process. Since this is a hard limit, the process that requires

more memory will page and appear to be performing poorly, even though the system, as a whole and other processes in the partition, performs well. Do not use Per Process Memory limits for a mix of processes that have widely varying memory requirements.

RPM causes other applications to fail or receive access violations

Cause: Memory limits for resource partitions are stringent, and applications are not allowed to break those limits. Thus, if an application asks for memory beyond the partition's memory limit, the application's request for memory will fail. An application can get an access violation as a result of using nonexistent memory.

Problem: Each application should check that memory allocation requests are successful. If the application in question does not perform such checks, the application might fail when the limit is reached. If the intention is to control a misbehaving application (such as a memory leaking application), then RPM is meeting its intended purpose. Restart the application normally if possible. If normal restarting procedures are not possible due to a suspended application, it may be possible to restart the application by stopping the partition and restarting the partition. If the application is being suspended too frequently, consider raising the memory limit for the partition.

RPM does not allow processes to run

Cause: RPM is applying a rule that may be too restrictive.

Solution: In a busy system, processes or partitions that belong to a lower priority or scheduling class can be deprived of memory. If there is not enough time for a lower priority process to run, it will either have to compete on an equal footing by having its priority raised (or other priorities lowered), or processors will have to be apportioned so the lower priority work gets some attention. For example, with eight processors, one processor might be set aside for low priority groups, while the higher priority groups compete to use the other seven processors. The low priority work will get some time and cannot be completely starved by the other work. If the starved process or group is run infrequently (or spends considerable time waiting), it might be desirable to place it at a higher priority than other work. Its work will then get carried out sooner, other work will be minimally affected, and overall throughput will increase.

Additional information

For additional information, refer to the Workload Management Pack Web site. This section contains Q&As that should be consulted before contacting HP Technical Support. This information is located at www.hp.com/products/wmp.

For more information

www.hp.com/go/proliant

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2002 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

All other product names mentioned herein may be trademarks of their respective companies.

TC020401BP, 03/2002

Printed in the US

